Restoration Schemes with Differentiated Reliability

Kai Wu, Luca Valcarenghi, and Andrea Fumagalli
Optical Networking Advanced Research (OpNeAR) Laboratory
Erik Jonsson School of Engineering and Computer Science
The University of Texas at Dallas
http://opnear.utdallas.edu
`kxw016500@utdallas.edu,{lucav,andreaf}@utdallas.edu`

***Abstract*—Reliability of data exchange is becoming increasingly important. In addition, applications may require multiple degrees of reliability. The concept of Differentiated Reliability (DiR) was recently introduced in [1] to provide multiple degrees of reliability in protection schemes that provision spare resources.**

With this paper, the authors extend the DiR concept to restoration schemes in which network resources for a disrupted connection along secondary paths are sought upon failure occurrence, i.e., they are not provisioned before the fault. The DiR concept is applied in two dimensions: restoration blocking probability — i.e., the probability that the disrupted connection is not recovered due to lack of network resources — and restoration time — i.e., the time necessary to complete the connection recovery procedure. Differentiation in the two dimensions is accomplished by proposing three preemption policies that allow high priority connections to preempt resources allocated to low priority connections. The three policies trade complexity, i.e., number of preempted connections, for better reliability differentiation.

Obtained results indicate that by using the proposed preemption policies, it is possible to guarantee a significant differentiation of both restoration blocking probability and restoration time. By carefully choosing the preemption policy, the desired reliability degree can be obtained, while minimizing the number of preempted connections.

I. INTRODUCTION

Emerging networked applications are characterized by Quality of Service (QoS) requirements such as bandwidth, delay, and reliability. In recent years, reliability is gaining an increasing importance because of service provider's need to satisfy the Service Level Agreement (SLA) offered to end users [2]. Reliability schemes must be implemented to mitigate or prevent service outages upon network failures.

In dynamic network scenarios, reliability can be assured by either dynamic protection schemes or restoration schemes. Dynamic protection schemes provision spare resources to protect connections upon their set up. Restoration schemes look for unreserved network spare resources along which failed connections reroute only upon failure occurrence. Due to this nature, restoration schemes are more applicable to highly dynamic networks with distributed control.

Connections sharing the same network may require different reliability degrees. For example mission-critical applications, such as medical and emergency-related applications and banking institution transactions may coexist with non-critical applications such as Internet radio, file transfer, web browsing, and e-mail. In [1], [3] the concept of *Differentiated Reliability (DiR)* is introduced for the design of static protection-based networks. According to the DiR concept each connection at the network layer under consideration is guaranteed a minimum absolute reliability degree, defined as the *Maximum Failure Probability (MFP)*, allowed for that connection. The reliability degree chosen for a given connection is determined by the application's QoS requirement. Dynamic protection schemes for providing dynamic connections with absolute differentiated reliability degrees have also been proposed [4].

In this study the authors extend the DiR concept to be applied to restoration schemes as follows. Each connection for which restoration is sought — termed as *restorable connection* — is characterized by a two-tuple: the *restoration blocking probability* defined as the probability of unsuccessful completion of the connection recovery procedure due to lack of resources, and the *restoration time* defined as the time required to complete the recovery procedure (successfully or not). Restorable connections can thus be differentiated into classes using the above two dimensions. High class connections must experience a lower restoration blocking probability and a shorter restoration time than low class connection.

Differentiation in both restoration blocking probability and restoration time is accomplished utilizing *preemption policies* [5]. Preemption policies have been proposed for providing available and reliable services to high-priorities connections in case of heavily loaded networks, or in case of link or node failures [5], [6], [7]. However, the ability of these preemption policies to differentiate the degree of reliability offered to various network connections is still to be assessed. Three preemption policies are studied that provide tangible differentiation of reliability: *restoration preemption* (RP), *working preemption* (WP), and *restoration and working preemption* (RWP).Under these policies, restoration attempts of highly reliable connections may preempt resources carrying working or restoration traffic of less reliable connections[1].

The simulation-based quantitative analysis presented in Section III-C shows that the three preemption policies are able to differentiate both the restoration blocking probability and the restoration time of connections that belong to distinct reliabil-

---

This work was supported in part by the NSF under contract #ANI-0082085.

Dr. Valcarenghi is currently with Scuola Superiore S. Anna, Pisa, Italy. E-mail: valcarenghi@sssup.it

[1] It is likely that after a network failure a number of connections are concurrently disrupted and need to be simultaneously restored.

ity classes. In addition, by carefully choosing the preemption policy to be used, the number of preempted connections upon network failure can be minimized without any significant loss of reliability differentiation.

II. DiR IN RESTORATION SCHEMES

Restoration schemes are utilized when connection demands are highly dynamic and, therefore, optimal provisioning of spare resources is not possible. Upon request, connections are created in the network reserving only working resources, i.e, those required to carry the connection without any spare capacity. The occurrence of a network failure, for example a link failure, triggers the rerouting of the affected connections, using available resources in the network. In the presence of light load, disrupted connections are likely restored promptly due to abundance of network resources. At medium and high loads, however, finding spare resources may be time consuming, and may not even be possible due to a lack of sufficient network resources. The DiR concept is introduced to provide differentiation of reliability at any load as discussed next.

A. Model Definition

The network under consideration is connection-oriented, e.g., a GMPLS based network. It is assumed that each established connection requires a dedicated and fixed amount of capacity, referred to as a channel. In a Wavelength Division Multiplexing (WDM) network, for example, the channel represents a wavelength. Each link can carry a number of connections equal to the number of available channels. The following definitions are used to formally introduce the DiR concept in restoration schemes.

- $\mathcal{C} = \{c_0, c_1, \cdots, c_m\}$, set of *reliability classes*;
- $\mathcal{C}_{lp}(c_i) = \{c_k : 0 \leq k < i\}$, subset of lower priority classes. Connections of class $c_k \in \mathcal{C}_{lp}$ require a lower reliability than connections belonging to class c_i;
- $G(\mathcal{N}, \mathcal{L})$, graph representing a network with a node set $\mathcal{N}$ and a bidirectional link set $\mathcal{L}$;
- $p_{s,d} = \{l_0, \cdots, l_k\}$, path between node pair (s, d) spanning links $l_k \in \mathcal{L}$;
- $d^{c_i}_{s,d}$, connection between node pair (s, d) belonging to reliability class c_i;
- $\mathcal{P}^{c_i}_{s,d}$, set of $kp^{c_i}_{s,d} = |\mathcal{P}^{c_i}_{s,d}|$ primary paths available for connections belonging to class c_i between node pair (s, d);
- $\mathcal{B}^{c_i}_{s,d}$, set of $kb^{c_i}_{s,d} = |\mathcal{B}^{c_i}_{s,d}|$ backup paths available upon failure for connections belonging to class c_i between node pair (s, d);
- β_l, total number of channels on link l;
- $\mathcal{E}^{c_i}_l$, set of $\mu p^{c_i}_l = |\mathcal{E}^{c_i}_l|$ channels reserved on link l by primary routes of connections belonging to class c_i;
- $\mathcal{F}^{c_i}_l$, set of $\mu b^{c_i}_l = |\mathcal{F}^{c_i}_l|$ channels reserved on link l by backup routes of connections belonging to class c_i;
- $\mathcal{G}^{c_i}_l$, set of $\eta p^{c_i}_l = |\mathcal{G}^{c_i}_l|$ primary connections belonging to class c_i disrupted because of the failure of link $l \in \mathcal{L}$;
- $\mathcal{H}^{c_i}_l$, set of $\eta^{c_i}_l = |\mathcal{H}^{c_i}_l|$ primary connections belonging to class c_i disrupted because of preemption after the failure of link $l \in \mathcal{L}$.

B. Protocol Description

In this section the protocol devised to implement the DiR concept in restoration schemes is described, with special attention paid to the restoration procedure. The term *path* is used to identify the set of links along which a connection may be routed. The term *route* is used to identify the actual path chosen to set up the connection.

Upon arrival at a network node, a request for a restorable connection is assigned a *reliability class* c_i. The connection is then routed along the shortest path in the set $\mathcal{P}^{c_i}_{s,d}$ with available channels. If no primary path $p_{s,d} \in \mathcal{P}^{c_i}_{s,d}$ has at least one channel available along each of its links the connection request is rejected. When set up, the connection is assigned a channel along the chosen primary route. At set up time of the connection, no spare resources are reserved. Upon network failure occurrence, e.g., a fiber cut, all the disrupted restorable connections dynamically look for spare channels along the pre-computed backup paths $p_{s,d} \in \mathcal{B}^{c_i}_{s,d}$.

Connection reliability differentiation is obtained by means of preemption policies that are described as follows.

1) Preemption Policies: As several connections that are disrupted by the network fault may look for spare resources concurrently, contention may occur during the restoration phase. Preemption policies are used in this case to resolve such contentions and provide multiple classes of reliability. The proposed three preemption policies are implemented in a distributed fashion and deployed at each network node.

The *restoration preemption* (RP), in which restoration attempts performed by high class connections can preempt channels previously reserved by backup routes chosen by low class connections. The preempted connections are therefore forced to choose an alternative backup route.

In the *working preemption* (WP) policy, restoration attempts of high class connections may preempt channels already reserved for primary routes of low class connections. In this case, connections that are not directly disrupted by the fault, may be indirectly disrupted by preemption. Therefore preempted primary connections need to activate the restoration procedure to find a backup route.

The *restoration and working preemption* (RWP) policy results from the combination of the previous two. In RWP, the choice of preemption policy adopted by a high class connection (either RP or WP) depends on a network-wide probabilistic parameter, p. When resource contention occurs, the restoration attempt of high class connections first try to preempt channels reserved by backup routes of low class connections utilizing RP. If the RP attempt fails, a WP attempt is made with a probability p, to preempt channels reserved for primary routes of low class connections. In particular, when $p = 0$, RWP reduces to RP; when $p = 1$, the preemption policy utilized is called *pure RWP*. The objective of this policy is to offer the required reliability differentiation while minimizing the number of primary connections preempted. Indeed WP is utilized only when RP fails.

In all policies, if a disrupted connection cannot successfully exert preemption, the restoration attempt is considered blocked.

2) Restoration Protocol: Restoration attempts are executed in a distributed manner. Network nodes are required to imple-

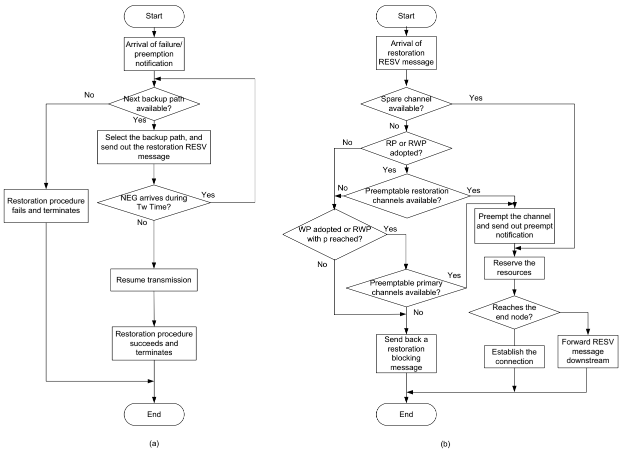

Fig. 1. Flowcharts of restoration protocol: a) Connection Master Node (CMN) functions; b) Link Master Node (LMN) functions.

ment two functions. First, a node is the *link master node* (LMN) for its outgoing links. The LMN is responsible for managing the associated link channels and the chosen preemption policy (see Fig. 1b). Second, the end-node of a disrupted connection closer to the faulty link is said to be the *connection master node* (CMN). The CMN is responsible for managing the connection restoration procedure (see Fig. 1a).

On receiving the failure notification of a disrupted connection $d_{s,d}^{c_i}$, the CMN chooses the shortest restoration path $p_{s,d}$ among the set of pre-computed backup paths $\mathcal{B}_{s,d}^{c_i}$ and sends a reservation message (RESV) along $p_{s,d}$. The RESV contains the set of links and nodes of $p_{s,d}$ and the class information of the connection. After sending out the RESV, the CMN waits a predefined time T_w for the negative acknowledgment messages (NEG) from LMNs along the chosen restoration path. If no NEG is received during T_w, the CMN resumes transmission and successfully completes the restoration procedure; otherwise, the CMN makes a successive restoration attempt along the next shortest unselected path in $\mathcal{B}_{s,d}^{c_i}$, and when all the paths in the set $\mathcal{B}_{s,d}^{c_i}$ have been tried, the restoration procedure terminates unsuccessfully.

Upon arrival of a RESV message, the LMN checks if a channel is available along the outgoing link l that belongs to the chosen backup path $p_{s,d}$ contained in the RESV. If at least one channel is available (i.e., $\beta_l - \mu_l > 0$, where $\mu_l = \sum_{i=0}^{m}(\mu p_l^{c_i} + \mu b_l^{c_i})$), the channel is reserved and the RESV is forwarded to the downstream node along $p_{s,d}$. If no spare channel is available (i.e., $\beta_l - \mu_l = 0$), the LMN activates the implemented preemption policy.

If RP is implemented, the set of channels reserved by backup routes of lower class connections, if not empty, (i.e., $\mathcal{F} = \{f | f \in \mathcal{F}_l^{c_j} \wedge c_j \in C_{lp}(c_i)\}$) can be preempted. If WP is implemented, the set of channels utilized by primary routes of lower class connections, if not empty, can be preempted (i.e., $\mathcal{E} = \{e | e \in \mathcal{E}_l^{c_j} \wedge c_j \in C_{lp}(c_i)\}$). If RWP is implemented, WP is tried with probability p if RP is tried first but fails. In all the preemption policies, the channel preempted is the one belonging to the lowest possible reliability class. If there are more than one such channels, the one that was most recently reserved is chosen.

If preemption is successful, the LMN forwards the RESV to the downstream node along $p_{s,d}$, and sends a preemption notification message upstream along the preempted route (backup route in RP and primary route in WP) to the CMN. If, instead, the preemption is unsuccessful, because the set of channels preemptable by the implemented preemption policy is empty, the LMN sends a restoration blocking message back to the CMN and resources previously reserved are released.

Both the restoration blocking message and the restoration preemption notification received by CMN are considered as NEG that triggers a successive restoration attempt. The working preemption notification message, however, is treated by CMN as a connection failure notification that triggers a new restoration of the preempted connection.

III. PERFORMANCE EVALUATION

In the first part of this section the performance evaluation metrics of interest are defined. Then the simulation scenario utilized is described. Finally, the results obtained are presented and discussed.

A. Performance Metrics

The three performance evaluation metrics utilized are the *restoration blocking probability (RBP)*, the *restoration time (RT)*, and the *failure propagation ratio (FPR)*. These metrics are evaluated independently for each of the connection classes in the network.

1) Restoration Blocking Probability: Under the assumption of single bidirectional link failure, the restoration blocking probability $RBP_l^{c_i}$ of the class c_i is defined as the ratio between the number of unsuccessfully restored connections B_l and the number of primary connections disrupted by both the failure of link $l \in \mathcal{L}$ and the preemption of higher class connection restoration attempts

$$RBP_l^{c_i} = \frac{B_l}{\mu p_l^{c_i} + \eta_l^{c_i}}, \tag{1}$$

$RBP_l^{c_i}$ depends on the class c_i, the number of attempts allowed for each disrupted connection $d_{s,d}^{c_i}$ (i.e., the number of pre-computed restoration paths available $kb_{s,d}^{c_i}$), the distribution of the connections to the different classes, and the preemption policy utilized.

2) Restoration Time: For a disrupted connection $d_{s,d}^{c_i}$ the restoration time $RT(d_{s,d}^{c_i})$ is the time elapsed between the occurrence of network failure and the termination, either successful or unsuccessful, of the connection restoration procedure. The restoration time $RT_l^{c_i}$ of a class c_i is defined as the average restoration time among all the primary connections belonging to reliability class c_i disrupted either by the failure of link $l \in \mathcal{L}$ or by the preemption of higher reliability class

$$RT_l^{c_i} = \frac{\sum_{d_{s,d}^{c_i} \in (\mathcal{G}_l^{c_i} \bigcup \mathcal{H}_l^{c_i})} RT(d_{s,d}^{c_i})}{\mu p_l^{c_i} + \eta_l^{c_i}}, \tag{2}$$

3) Failure Propagation Ratio: The failure propagation ratio of a class c_i is the ratio between the preempted primary connections and the connection disrupted by the failure of link l.

$$FPR_l^{c_i} = \frac{\eta_l^{c_i}}{\mu p_l^{c_i}}, \qquad (3)$$

This metrics measures the number of additional primary connection disruptions caused by working preemption.

With the assumption that the probability of single link failure is uniformly distributed among all the network links, the expected values for all the former three metrics are defined as

$$E\{M^{c_i}\} = \frac{\sum_{l=0}^{|\mathcal{L}|-1} M_l^{c_i}}{|\mathcal{L}|}. \qquad (4)$$

where M^{c_i} is the metric considered for the class c_i.

B. Simulation Setup

A customized $C++$ based event driven simulator is utilized to evaluate the performance of the proposed implementation of the DiR concept in restoration schemes. Each bidirectional network link $l \in \mathcal{L}$ of the network $G(\mathcal{N}, \mathcal{L})$ is assigned a capacity of β_l channels. Each node is associated with an independent traffic generator which generates connection requests with Poisson distribution at a mean arrival rate of λ. The destination node is uniformly chosen among the remaining nodes. The connection class is also uniformly chosen among a set of m reliability classes. The connection duration time is exponentially distributed with unit mean. Thus the traffic intensity generated by each node is λ measured in Erlang.

The restoration time is measured in terms of hops spanned by the restoration messages. It is assumed that the time a message takes to propagate from one LMN to the adjacent LMN and to be processed at the receiving LMN is exactly one hop unit time. T_w is hence predefined as the hop count of the corresponding backup path plus two hop units, considering the computation delay at CMN and the establishment delay of the whole backup route.

Simulation is performed in two phases: the pre-failure phase and the post-failure phase.

1) Pre-Failure Phase: In this phase, requested connections are routed along one of the available primary paths. In the considered simulation scenario, connections between the same (s, d) pair belonging to different classes utilize the same set of primary and backup paths, i.e., $\mathcal{P}_{s,d}^{c_i} = \mathcal{P}_{s,d}$ and $\mathcal{B}_{s,d}^{c_i} = \mathcal{B}_{s,d}, \forall i$. Given a node pair (s, d) the sets of pre-computed paths are calculated as follows:

a) $\mathcal{P}_{s,d}$ is the set of k shortest link-disjoint paths between s and d [8].

b) For each primary path $p_{s,d}^j$, a set $\mathcal{B}_{s,d}(p_{s,d}^j)$ of secondary paths is calculated by applying the Martin's K-Ranking algorithm [9] to the graph $G(\mathcal{N}, \mathcal{L}^{'})$ obtained from the graph $G(\mathcal{N}, \mathcal{L})$ by eliminating the links $l \in p_{s,d}^j$.

The computed paths are sorted by increasing number of hops.

2) Post-Failure Phase: After the network reaches a steady state under a given node traffic intensity, the connection request arrival process is stopped (i.e., the network status is frozen), and a single link failure is triggered[2]. The restoration protocol described in section II-B.2 is then activated to recover the connections affected by the failure. After all the restoration procedures terminate, the network resources are refreshed back to the status at which the network was frozen, and another link failure event is triggered. The procedure repeats until all single link failure scenarios have been analyzed.

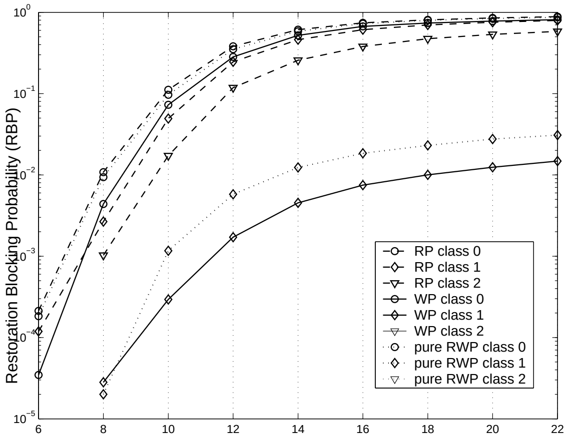

Fig. 2. Expected RBP for various preemption policies.

C. Results

This section presents the results obtained by simulations run on the NSF network [10]. The network topology consists of $|\mathcal{N}| = 16$ nodes and $|\mathcal{L}| = 25$ bidirectional links. Each link l has a capacity of $\beta_l = 32$ channels and the number of connection classes in the network is $|\mathcal{C}| = 3$. The maximum number of primary and backup paths are respectively $kp_{s,d} = 3$ and $kb_{s,d} = 5$ for each node pair (s, d). The number of links (hops) spanned by each path is bounded by $H = 6$. A number of 1280 experiments is run for each λ ranging from 6 to 22 Erlangs. Each simulation achieves a confidence interval of 5% or less at the 95% confidence level.

Fig. 2 and Fig. 3 show that both, RP and WP, preemption policies are able to differentiate both RBP and RT [3] among connections belonging to different reliability classes. However RP is not able to well differentiate between class 1 and class 0 connections. On the other hand, WP is not able to differentiate clearly the RBP and the RT of class 2 and class 1 connections. In addition, as shown in Fig. 6 WP yields high FPR. The introduction of RWP permits to tune the reliability differentiation of connections belonging to different classes, both in terms of RBP and RT. As shown in Fig. 4 and Fig. 5, by varying the value of p the relative differentiation between different reliability class connections changes. In addition by choosing an appropriate value for p, it is possible, as shown in Fig. 6, Fig. 4,

---

[2] The simulation freezes the network to allow separate operations of restoration procedure and primary connection add/drop, under the assumption that these two operations take two different time scales and rarely overlap in time.

[3] The RT for class 0 connections drops slightly as load increases, because the interval between successive attempts of class 0 connections shortens more than the increasing number of attempts due to the increasing blocking and preemption probability.

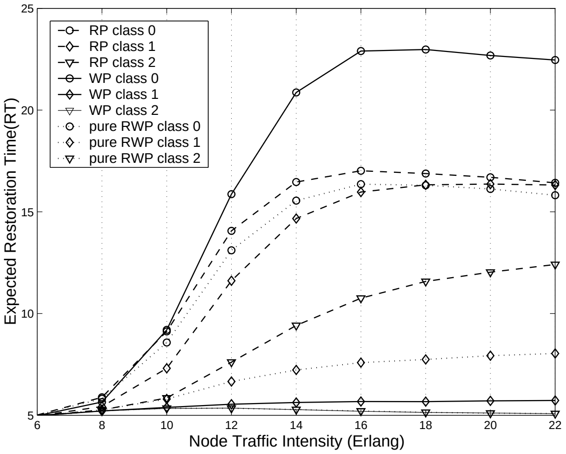

Fig. 3. Expected RT for various preemption policies.

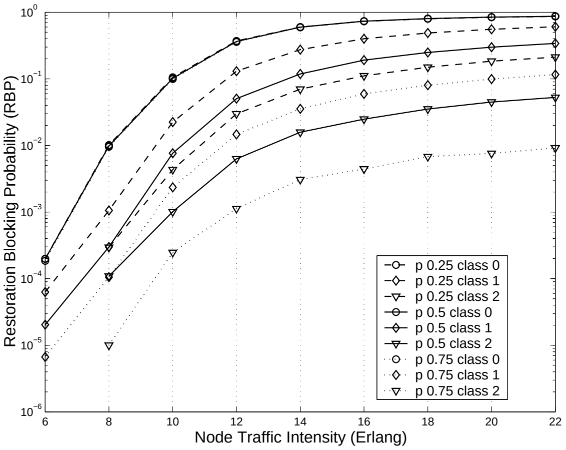

Fig. 4. Expected RBP for RWP with different values of p.

and Fig. 5, to minimize the FPR for the lowest priority class while maintaining reliability differentiation.

IV. Conclusion

The paper extended the concept of Differentiated Reliability (DiR) to restoration schemes. Connections belonging to high reliability classes are guaranteed lower Restoration Blocking Probability (RBP) and faster Restoration Time (RT) than the ones experienced by connections belonging to low reliability classes. Such differentiation is accomplished by means of preemption policies that facilitate restoration attempts of disrupted connections belonging to high reliability classes.

Results show that the proposed three preemption policies, namely Restoration Preemption (RP), Working Preemption (WP) and Restoration and Working Preemption (RWP), are able to guarantee differentiated classes of reliability in terms of both RBP and RT. In particular, the RWP policy permits, through the use of a probabilistic approach, to fine tune the reliability differentiation. By carefully choosing the probability of preempting primary connections, it is possible to decrease the number of preempted primary connections, while at the same time maintaining tangible reliability differentiation.

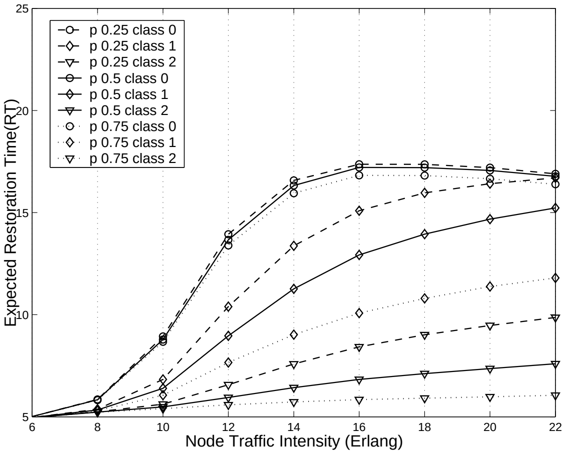

Fig. 5. Expected RT for RWP with different values of p.

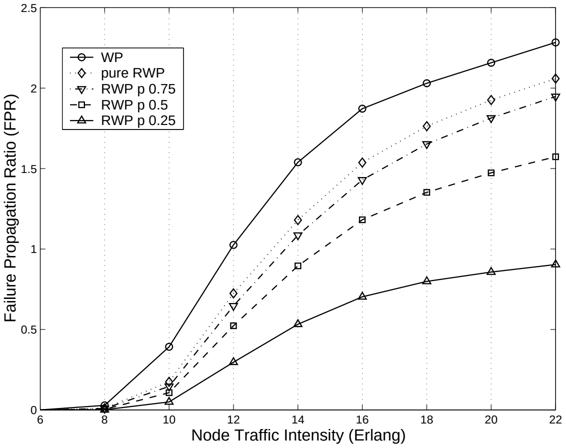

Fig. 6. Expected FPR for class 0 connections.

References

[1] A. Fumagalli and M. Tacca, "Optimal Design of Differentiated Reliability (DiR) Optical Ring Networks," in *International Workshop on QoS in Multiservice IP Networks (QoS-IP) 2001*, January 2001.

[2] E. Marilly, O. Martinot, H. Papini, and D. Goderis, "Service Level Agreements: A Main Challenge for Next Generation Networks," in *2nd European Conference on Universal Multiservice Networks, EDUMN 2002*, 2002, pp. 297–304.

[3] O. Gerstel and G. Sasaki, "Quality of Protection (QoP): A Quantitative Unifying Paradigm to Protection Service Grades," *Optical Networks Magazine*, vol. 3, no. 3, pp. 40–49, May/June 2002.

[4] C. V. Saradhi and C. S. R. Murthy, "Routing Differentiated Reliable Connections in WDM Optical Networks," *Optical Networks Magazine*, vol. 3, no. 3, pp. 50–67, May/June 2002.

[5] J.A. Garay and I.S. Gopal, "Call Preemption in Communication Networks," in *INFOCOM '92*, 1992, vol. 3, pp. 1043–1050.

[6] M. Peyravian and A.D. Kshemkalyani, "Connection Preemption: Issues, Algorithms, and a Simulation Study," in *INFOCOM '97*, 1997, vol. 1, pp. 143–151.

[7] J. C. de Oliveira, C. Scoglio, I. F. Akyildiz, and G. Uhl, "A New Preemption Policy for DiffServ-Aware Traffic Engineering to Minimize Rerouting," in *INFOCOM 2002*, 2002, vol. 2.

[8] R. Bhandari, *Survivable Networks — Algorithms for Diverse Routing*, Kluwer Academic Publishers, 1999.

[9] E. Q. Vieira Martins, M. M. B. Pascoal, and J. L. E. Santos, "An Algorithm for Ranking Loopless Paths," Tech. Rep., CISUC, December 1999.

[10] H. Zang, J. P. Jue, L. Sahasrabuddhe, R. Ramamurthy, and B. Mukherjee, "Dynamic Lightpath Establishment in Wavelength-Routed WDM Networks," *IEEE Communications Magazine*, vol. 39, no. 9, pp. 100–108, Sep. 2001.